

IN THE SPECIFICATION:

Please revise the specification as follows:

Please amend paragraph numbers [0038], [0039], [0052], and [0053] as follows:

**[0038]** Memory 218 preferably includes high-speed random access memory and may include non-volatile memory, such as one or more magnetic disk storage devices. In some implementations, ~~the~~ memory 218 may include mass storage that is remotely located from the central processing unit(s) 204, such as the shared storage 108 shown in Figure 1. ~~The memory~~ Memory 218 preferably stores an operating system 220, such as Linux, Unix or Windows (a trademark of Microsoft), that includes procedures for handling basic system services and for performing hardware dependent tasks. ~~Memory 218, or alternatively one or more storage devices (e.g., one or more nonvolatile storage devices) within memory 218, includes a computer readable storage medium.~~ ~~Memory 218, or computer readable storage medium,~~ also preferably stores communications procedures 222 used for communicating with other devices on the network 100 (Figure 1), such as other slaves, the master, and/or the generator. In particular, the communication procedures 222 are used for transmitting the slave state to the master, receiving instructions from the master, requesting and/or receiving data from other slaves or the generator, as described below in relation to Figure 4.

**[0039]** ~~The memory~~ Memory 218 also preferably includes: a file set file (FileSet) 224; the slave state 226; data files 228; checksum files 230; verification procedures 240; failure detection procedures 242; state determination and reporting procedures 244; and a cache 246 for temporarily storing data. The file set file (FileSet) 224 contains a list of files required by the particular slave. In some embodiments, there are multiple distinct sets of slaves, each of which has a different set of files. For instance, if the system 100 (Figure 1) is a search engine system, there may be anywhere from five to one hundred (or perhaps a few hundred) distinct types or classes of slaves, each of which performs a distinct function or set of functions in the process of responding to search engine queries. For instance, the index of documents to be searched may be divided over a small or large number of slaves, and therefore each of these slaves will have a distinct FileSet. Furthermore, one or more of the datacenters of the system may have more than one copy of each type of slave, thereby providing multiple replicas of the entire search engine so as to be able to service large numbers of search requests. Thus, if the system has slaves of types 1 through N, and a datacenter has M replicas of the search engine, where M is an integer greater than 1, there will be M slaves of each type in the data center. These slaves will preferably be spread over

multiple racks so that no single rack failure will prevent the system from being able to respond to search queries.

[0052] Memory 362 preferably includes an operating system (OS) 376, such as Linux, Unix or Windows (a trademark of Microsoft) that includes procedures for handling basic system services and for performing hardware dependent tasks. Memory 362, or alternatively one or more storage devices (e.g., one or more nonvolatile storage devices) within memory 362, includes a computer readable storage medium. Memory 362, or computer readable storage medium, also preferably includes communications procedures 378 that are generally used for communicating with the network 100 (Figure 1) and slaves. In particular, the communication procedures 378 are used for receiving a slave's state and for transmitting instructions to slaves, as described below in relation to Figure 4.

[0053] The memory Memory 362 also preferably includes a system hierarchy table 380, a system resources table 382, a pending transmissions table 384, a state table 390, failure determination procedures 392, resource allocation and scheduling procedures 394, and a cache 396. The system hierarchy table 380 is used by the master to determine the proximity of slaves to one another. The system resources table 382 is used by the master to determine the bandwidth resources for each communication path in the switched network, and the resources currently in use or reserved for use. The pending transmissions table 384 is used to keep track of all current or pending transmissions. The state table 390 is used by the master to determine which system resources have failed, as well as which files or blocks of files have been received by each slave and while files or file blocks are still needed by each slave. The failure determination procedures 392 are used by the master for determining whether a resource has failed. The resource allocation and scheduling procedures 394 are used to schedule data transfers between slaves, as described in further detail below in relation to Figures 4 and 5A-5D. The cache 396 is used for temporarily storing data.